Shield and Blindfold: Agentic AI, Anonymity, and the Civil Rights Inversion

Anirban Mukherjee
Hannah Hanwen Chang

Draft — April 20, 2026

Anirban Mukherjee (anirban@avyayamholdings.com) is Principal at Avyayam Holdings. Hannah H. Chang (hannahchang@smu.edu.sg; corresponding author) is Associate Professor of Marketing at the Lee Kong Chian School of Business, Singapore Management University. This research was supported by the Ministry of Education (MOE), Singapore, under its Academic Research Fund (AcRF) Tier 2 Grant, No. MOE-T2EP40124-0005.

Abstract

AI agents can transact with retailers and platforms on behalf of a human principal without revealing whom they represent, returning to commerce a degree of anonymity that has not been practically available since cash. The civil rights implications run in both directions. The architecture of discrimination collapses when the buyer is a shell—a retailer that cannot identify its counterparty cannot proxy-discriminate. But the same collapse defeats the evidentiary infrastructure on which civil rights enforcement relies. Disparate-treatment doctrine needs comparators. Disparate-impact doctrine needs a protected class within the defendant's customer base. *Iqbal*-era pleading rigor needs facts the anonymity regime prevents from existing. The mechanism that forecloses discrimination also forecloses its proof. This Article argues that the resulting civil rights inversion requires the law to decide whether agent-mediated anonymity should be treated as civil rights infrastructure, how abuse can be regulated without forcing reidentification, and whether retailers may refuse to deal with agents at all.

Keywords: Privacy, Civil Rights, Artificial Intelligence, Antidiscrimination Law, Algorithmic Discrimination.

JEL codes: K38, K24, O33, L86.

Table of Contents

Introduction

> "Nothing was your own except the few cubic centimetres inside your skull."
>
> — George Orwell, *Nineteen Eighty-Four* (1949)

Each technology that reshapes modern life arrives with the same dread: that it will hollow out the private sphere.[1] The printing press prompted fears that ideas would be surveilled at the point of publication.[2] The telegraph and the telephone put voices into wires that, it was imagined, someone was always listening to.[3] The industrial city was an alienating nightmare to those who remembered the village.[4] The credit card replaced the unmarked handshake of cash

[1] *See generally* Daniel J. Solove, Understanding Privacy 4–11 (2008) (tracing recurring privacy anxieties across successive information technologies).

[2] *See* Elizabeth L. Eisenstein, The Printing Press as an Agent of Change 344, 347–48 (1979) (describing the Archbishop of Mainz's 1485 licensing edict and subsequent papal censorship decrees prompted by the printing press).

[3] *See* Olmstead v. United States, 277 U.S. 438, 473–74 (1928) (Brandeis, J., dissenting) (warning that government wiretapping could obtain "what is whispered in the closet"); Tom Standage, The Victorian Internet 110–11 (1998) (describing widespread anxiety that telegraph operators and intermediaries could read every private message in transit).

[4] *See* Raymond Williams, The Country and the City 1–12 (1973) (tracing the recurring literary contrast between an idealized rural past and a corrupt, alienating urban present).

with a permanent, centralized ledger of every purchase.[5] Orwell's image of the last redoubt—those few cubic centimetres inside the skull—is the most elegant expression of an anxiety that is, by now, several centuries old.[6] In each iteration, the fear is the same: the new machine sees too much, remembers too well, and shares too widely.[7]

Artificial intelligence, however, contains its own inversion. As cryptocurrency shows, the digital age need not end anonymous exchange.[8] AI agents extend this principle to the transaction as a whole: an AI agent can serve as an ephemeral intermediary, indistinguishable from any other retail-platform user, that browses, pays, receives, and reviews for a human principal without revealing whom it represents.[9] It can pay in cryptocurrency or a single-use card, receive goods at an anonymized address, and leave no readily accessible record linking the transaction to the principal at the point of sale.[10] Anonymity in commerce can be restored—not by retreating from technology, but through it.

I. Civil Rights Starved

Commercial civil rights law rests on an assumption so foundational it is rarely stated: the parties to a transaction are identifiable.[11] Discrimination, as the law understands it, is something a seller

[5] *See* Josh Lauer, Creditworthy: A History of Consumer Surveillance and Financial Identity in America 235–63 (2017).

[6] George Orwell, Nineteen Eighty-Four pt. I, ch. 2 (1949) ("Nothing was your own except the few cubic centimetres inside your skull.").

[7] *See* Solove, *supra* note 1.

[8] Bitcoin and its successors record transactions on a public ledger without necessarily linking them to real-world identities. *See* Satoshi Nakamoto, *Bitcoin: A Peer-to-Peer Electronic Cash System* 1, 6, 10 (Oct. 31, 2008) (unpublished manuscript), https://bitcoin.org/bitcoin.pdf (describing a system where "the public can see that someone is sending an amount to someone else, but without information linking the transaction to anyone"); Primavera De Filippi & Aaron Wright, Blockchain and the Law: The Rule of Code 20–21 (2018).

[9] *See* Yonadav Shavit et al., *Practices for Governing Agentic AI Systems* 4–9 (OpenAI Research Paper, Dec. 14, 2023), https://openai.com/research/practices-for-governing-agentic-ai-systems (defining agentic AI systems as those that "pursue complex goals with limited direct supervision"); Noam Kolt, *Governing AI Agents*, 101 Notre Dame L. Rev. (forthcoming 2026) (manuscript at 8–12), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4772956 (proposing a governance framework for AI agents grounded in agency law principles).

[10] *See supra* note 8.

[11] *See* Samuel R. Bagenstos, *The Structural Turn and the Limits of Antidiscrimination Law*, 94 Cal. L. Rev. 1, 10–14 (2006) (analyzing antidiscrimination law's dependence on identifying perpetrators and victims of discriminatory

does to a buyer the seller has categorized. Remedy is something a court imposes after a plaintiff, a member of a protected class, shows she was treated differently from a similarly situated person outside it.[12] Both presuppose that the seller knows who is buying and that the law can prove it.

This assumption is benign in the traditional marketplace. Consider a reader who wishes to buy a book that carries some social or professional risk: a political memoir, a medical reference, a work on a stigmatized topic. In the ordinary online marketplace the purchase leaves a trail: the search query ties to an account, the order to a name and billing address, the payment to a credit card, the shipment to a home, and subsequent browsing surfaces similar titles across every site in the advertising ecosystem.[13] Each trace is individually modest; stitched together, they form a durable profile.

An AI agent, however, alters the picture. It can present as an ordinary user through an account the agent controls, not one tied to the principal's identity, pay with a single-use virtual card or prefunded cryptocurrency, and direct shipment to a parcel locker the principal retrieves without showing identification linked to the purchase.[14] The retailer sees a transaction indistinguishable from any other but cannot link that user to the principal or to the principal's other transactions.[15]

The civil rights consequences of AI are unusual precisely because they run in both directions. If a retailer cannot identify the individual behind a transaction, it cannot proxy-discriminate: it cannot price by zip code, steer by inferred demographics, target by protected class, for there is no data from which to infer when the counterparty is an anonymous shell.[16] The architecture of algorithmic discrimination collapses. Anonymity, long thought a threat to civil rights, becomes

---

conduct as a precondition for liability).

[12] *See infra* note 24.

[13] *See* Joseph Turow, The Daily You 4–5, 91–103 (2011).

[14] *See supra* notes 9, 8.

[15] Each of the agent's steps has an analogue already in wide use: privacy-preserving browsing for tracking avoidance, virtual cards for payment, parcel lockers for receipt, and temporary email for communication. *See, e.g., Tor Browser*, Tor Project, https://www.torproject.org (last visited Apr. 8, 2026); *Privacy — Seamless & Secure Online Card Payments*, Privacy.com, https://privacy.com (last visited Apr. 8, 2026); *Amazon Hub Locker*, Amazon, https://www.amazon.com/b?node=6442600011 (last visited Apr. 8, 2026); *SimpleLogin — Open Source Email Alias Solution*, Proton, https://simplelogin.io (last visited Apr. 8, 2026).

[16] *See infra* notes 19, 20, 21, and 22.

their architecture.

But the same collapse defeats civil rights law's remedial architecture. Antidiscrimination enforcement depends on showing that a protected class was treated differently, which presupposes identifying who was treated, by whom, and how.[17] Affirmative remedies depend on identifying their beneficiaries. The mechanism that forecloses discrimination forecloses its proof and its remedy.

A. Discrimination starved of its inputs

The discrimination that modern civil rights law is principally concerned with is no longer the blunt refusal of service practiced in the 1960s.[18] It is statistical, algorithmic, and often unintentional: price discrimination keyed to zip-code models,[19] product steering driven by personalization signals,[20] credit scoring that relies on proxies for race and age,[21] ad targeting built on behavioral aggregation.[22] Each mechanism consumes the same input, demographic or behavioral data tied to an identifiable buyer, producing discrimination as a byproduct. A retailer cannot steer products

[17] *See infra* note 24.

[18] *See* Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 Cal. L. Rev. 671, 673–77 (2016) (describing how contemporary discrimination arises from facially neutral data-mining practices rather than intentional exclusion); Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 Wm. & Mary L. Rev. 857, 860–61 (2017) (contrasting algorithmic "classification bias" with the overt exclusion targeted by mid-twentieth-century civil rights statutes); *cf.* Heart of Atlanta Motel, Inc. v. United States, 379 U.S. 241, 261 (1964) (upholding Title II of the Civil Rights Act of 1964 against the paradigmatic mid-century practice of categorical refusal to serve Black patrons).

[19] Jennifer Valentino-DeVries, Jeremy Singer-Vine & Ashkan Soltani, *Websites Vary Prices, Deals Based on Users' Information*, Wall St. J., Dec. 24, 2012, at A1 (documenting Staples.com's practice of varying prices based on a user's inferred location and proximity to competitor stores); Ryan Calo, *Digital Market Manipulation*, 82 Geo. Wash. L. Rev. 995, 999–1004 (2014) (theorizing personalized and geographically targeted pricing as a form of algorithmic consumer harm).

[20] *See, e.g.*, Anikó Hannák et al., *Measuring Price Discrimination and Steering on E-commerce Web Sites*, in Proc. 2014 ACM Internet Measurement Conf. 305, 308–12 (2014).

[21] *See, e.g.*, Robert Bartlett et al., *Consumer-Lending Discrimination in the FinTech Era*, 143 J. Fin. Econ. 30, 42–44 (2022) (finding that algorithmic lenders charge Latinx and Black borrowers higher interest rates than otherwise-similar white borrowers, even when no human loan officer is involved); Anya E.R. Prince & Daniel Schwarcz, *Proxy Discrimination in the Age of Artificial Intelligence and Big Data*, 105 Iowa L. Rev. 1257 (2020) (theorizing how AI systems trained on facially neutral data systematically reconstruct protected characteristics through correlated proxies).

[22] *See, e.g.*, Latanya Sweeney, *Discrimination in Online Ad Delivery*, 56 Comm. ACM 44 (2013); Muhammad Ali et al., *Discrimination Through Optimization: How Facebook's Ad Delivery Can Lead to Biased Outcomes*, 3 Proc. ACM Hum.-Comput. Interaction, art. 199, at 12–15 (Nov. 2019); Charge of Discrimination, *Sec'y, U.S. Dep't of Hous. & Urb. Dev. v. Facebook, Inc.*, FHEO No. 01-18-0323-8 (Mar. 28, 2019).

away from a protected class without knowing the buyer's class membership. A lender cannot price a loan by zip code without geolocating the applicant. An ad network cannot target by inferred race if the inference has nothing to land on.

When the counterparty is an AI agent, the inputs are gone. The agent does not disclose the principal's location, inferred demographics, purchase history, or protected class. Each session begins without history; each session ends without persistence. The agent is a shell: no stable identity to profile, no durable link to the human behind it. Proxy discrimination, a dominant modern form, collapses.[23] The retailer cannot engage in it. The algorithm has no signal to fit. The discrimination machinery, deliberate or inadvertent, is starved of its substrate.

B. Proof starved of its evidence

The same mechanism that defeats discrimination defeats the enforcement apparatus built to detect and punish it. Disparate-treatment doctrine requires the plaintiff to show she was treated worse than a similarly situated counterparty outside her protected class.[24] That showing requires identifying herself (the plaintiff), the defendant's knowledge (what the defendant could have known about her), and a comparator (someone similarly situated but outside the class, who was treated differently). Each step depends on records that AI agents do not produce. The plaintiff identifies herself only by stepping out of the anonymity regime and disclosing the information the agent was built to conceal; even then, the defendant may have no record of who she was during the transaction. The comparator is an agent, indistinguishable from the plaintiff's own agent by design. The defendant's knowledge is null because the transactional interface provides nothing to know.

Disparate-impact doctrine fares no better. *Griggs v. Duke Power Co.* and its progeny require proof that a facially neutral practice produces statistically significant differential effects on a

[23] *See supra* note 21.

[24] *See* McDonnell Douglas Corp. v. Green, 411 U.S. 792, 802–04 (1973) (establishing the burden-shifting framework for disparate-treatment claims).

protected class.[25] That proof requires identifying the protected class within the defendant's customer base and measuring outcomes against that identification. If every customer is an agent, there is no protected class to identify, no demographic baseline to compare against, and no statistical machinery to deploy. The doctrine does not fail because the practice is fair. It fails because the data to test fairness does not exist. Discovery does not save the case: records the defendant does not keep cannot be produced, and a retailer who sees only agents has no customer-level demographics to disgorge. Expert testimony encounters the same wall. An expert retained to show that a retailer's checkout flow produces differential outcomes across protected classes needs the ground truth of user identity; a retailer whose records contain only agents has no such ground truth to provide. The expert is left testifying about a pattern in data that, under an anonymity regime, does not exist.

C. Pleading starved of its facts

The problem with enforcement is not only that the evidentiary record is thin at trial. It is that the record does not exist at the threshold at which enforcement begins. *Iqbal* and *Twombly* elevated the pleading standard from "conceivable" to "plausible," requiring a complaint to allege concrete facts from which unlawful conduct can be reasonably inferred.[26] A plaintiff must come to the courthouse with facts sufficient to make discrimination a *plausible*, not merely conceivable, explanation for the defendant's conduct.

Those facts come from three sources. The plaintiff may have *direct evidence*: a memo, an email, a policy, a remark revealing bias.[27] She may have *comparative evidence*: knowledge that similarly

[25] Griggs v. Duke Power Co., 401 U.S. 424, 431 (1971); *see also* Tex. Dep't of Hous. & Cmty. Affairs v. Inclusive Cmtys. Project, Inc., 576 U.S. 519, 539–40 (2015) (affirming disparate-impact liability under the Fair Housing Act).

[26] Ashcroft v. Iqbal, 556 U.S. 662, 678 (2009); Bell Atl. Corp. v. Twombly, 550 U.S. 544, 570 (2007).

[27] *See* Price Waterhouse v. Hopkins, 490 U.S. 228, 251–52 (1989) (plurality opinion) (recognizing that a plaintiff may prove discrimination by showing that an illegitimate criterion "played a motivating part" in the employment decision, including through direct expressions of bias); Michael J. Zimmer, *The New Discrimination Law: Price Waterhouse Is Dead, Whither McDonnell Douglas?*, 53 Emory L.J. 1887, 1889–92 (2005) (tracing the doctrinal role of direct evidence in disparate-treatment cases and the development of the "stray remarks" doctrine in the lower courts).

situated counterparts outside her protected class were treated differently.[28] Or she may have *statistical evidence*: patterns in a defendant's behavior revealed by prior litigation, investigative journalism, regulatory disclosure, or organized scrutiny.[29] In a world of identifiable customers, at least one of these is usually available. The architecture of discrimination enforcement depends on that availability.

Agent-mediated anonymity removes all three sources at once. There is no direct evidence, because the agent strips the demographic signal that would trigger discriminatory behavior: no email or memo can reveal bias against a class the defendant never knew it was transacting with. There is no comparative evidence, because comparators are themselves anonymous agents, indistinguishable from the plaintiff's own agent by construction. And there is no statistical evidence, because no defendant, regulator, or third party can assemble demographic data about customers who present only as shells. The pleading standard, faithfully applied, produces dismissal, not because the plaintiff's claim is implausible, but because the facts that would make *any* such claim plausible cannot exist.

The result is a doctrinal closed loop. A plaintiff cannot plead what she cannot know; she cannot know what she cannot discover; she cannot discover without first pleading. The gate locks from both sides. *Iqbal* is not malfunctioning; the Court's concern about abusive pleading is met exactly as intended. But *Iqbal*-era pleading rigor combined with agent-mediated anonymity forecloses the courthouse door to discrimination claims previously colorable under the permissive pleading regime.[30] The enforcement apparatus is not merely weakened. Its entry point is closed.[31]

---

[28] *See McDonnell Douglas*, 411 U.S. at 804 (suggesting that a plaintiff may show pretext through evidence that "white employees involved in acts against [the employer] of comparable seriousness … were nevertheless retained or rehired"); Suzanne B. Goldberg, *Discrimination by Comparison*, 120 Yale L.J. 728, 732–35 (2011) (analyzing the centrality and limits of the "similarly situated" comparator requirement in antidiscrimination doctrine).

[29] *See* Hazelwood Sch. Dist. v. United States, 433 U.S. 299, 307–08 (1977) (reaffirming that "gross statistical disparities" can themselves constitute prima facie proof of a pattern or practice of discrimination); Int'l Bhd. of Teamsters v. United States, 431 U.S. 324, 339–40 (1977) (recognizing that "statistics … come in infinite variety" and may be probative of pattern-or-practice discrimination).

[30] *See* Conley v. Gibson, 355 U.S. 41, 45–46 (1957) (establishing the "no set of facts" pleading standard subsequently overruled by *Twombly*).

[31] The argument that the pleading stage, rather than the ultimate burden of proof, is the operative failure point for discrimination claims in an anonymity regime appears novel to AI. Prior literature on discrimination pleading after *Iqbal* has focused on the evidentiary difficulties faced by particular classes of plaintiffs in particular doctrinal settings,

The same closure affects the affirmative side of the equality project. Antidiscrimination law is not only about preventing unequal treatment in the moment; it also remedies the effects of historical marginalization through targeted intervention: affirmative action in contracting, minority business development programs, targeted lending initiatives.[32] Each presupposes that the beneficiary can be identified as a member of the target group. At the point of transaction, anonymity strips that identification: an agent does not announce its principal's race to a minority business development program; a cryptocurrency payment does not identify its source community; a minority business set-aside cannot operate when bidders are agents. At the transaction, the mechanism that prevents harm also prevents help.

II. The Role of Law

The anonymity regime invites three responses from the law: whether anonymity should be treated as civil rights infrastructure, available to those who cannot afford it; how to permit legitimate use while foreclosing abuse; and whether retailers may refuse to transact with agents at all—and if so, what that refusal means for the anonymity this Article has described.

A. Anonymity as civil rights infrastructure

Sophisticated AI agents are not free. They require technical literacy, computational resources, and a willingness to navigate a commercial world that has not yet adapted to them. Early adopters will skew toward the populations that already enjoy the best privacy tools: the wealthy, the technically fluent, and those with the time to configure their transactional lives carefully.

---

see, e.g., Joseph A. Seiner, *The Trouble with Twombly: A Proposed Pleading Standard for Employment Discrimination Cases*, 2009 U. Ill. L. Rev. 1011; Charles A. Sullivan, *Plausibility Pleading Employment Discrimination*, 52 Wm. & Mary L. Rev. 1613 (2011), without considering the combinatorial problem posed when the underlying identification architecture is absent altogether. The closed-loop structure described here parallels the "intermediate copying double bind" identified in the AI copyright context: claims cannot be initiated without the very evidence that the structure of the technology prevents from existing. *See* Anirban Mukherjee & Hannah Hanwen Chang, *Compactibility: AI and the Idea-Expression Inversion* (Feb. 21, 2026) (unpublished manuscript), https://ssrn.com/abstract=6232359./label%7Bfn:novelty}

[32] *See, e.g.,* Exec. Order No. 11,246, 3 C.F.R. 339 (1964–1965); 15 U.S.C. § 637(a) (2024); 12 U.S.C. § 2901 (2024).

This is the digital divide at a new layer. Agent-mediated commerce will leave non-agent users exposed to the price discrimination, behavioral targeting, and algorithmic harms that agent-users escape. The discrimination machinery does not disappear; it shifts from targeting members of a protected class to targeting the residual population that did not, or could not, deploy the shield. The defeat of discrimination against agent-users is not the defeat of discrimination; it is its concentration on those who remain visible.

The obvious response is to sever the link between anonymity and wealth. If privacy-enhancing agents become the most effective protection against the discrimination machinery Part I described, access to those agents should not be contingent on the ability to pay. Civil rights infrastructure has addressed analogous problems before. Legal Aid exists because legal representation is expensive; the EEOC exists because individual enforcement is resource-intensive; housing counselors exist because the information asymmetries in mortgage and rental markets are severe.[33] Each rests on the same principle: where a protection is necessary to prevent a recognized form of discrimination, access to the protection should not be limited to those who can afford it.

A parallel response would treat the AI agent itself as civil rights infrastructure—publicly provided, subsidized, or required to be offered free by platform providers above a certain scale. The precedents are imperfect but not absent. Public libraries provide free internet access because connectivity is now necessary for participation in civic life. Universal service rules subsidize telecommunications in underserved areas because the telephone network became too important to leave to unsubsidized markets.[34]

Agent-mediated anonymity could be framed in the same register, if the law recognizes that what is being protected is not merely consumer preference but the structural condition on which civil rights enforcement in a post-identification commercial environment depends. Public libraries and universal service are precedents for access to a positive good; the closer precedent is *Gideon v.*

[33] *See* Alan W. Houseman & Linda E. Perle, Securing Equal Justice for All: A Brief History of Civil Legal Assistance in the United States (Ctr. for L. & Soc. Pol'y rev. ed. 2018); 42 U.S.C. § 2000e-4 (2024) (establishing the Equal Employment Opportunity Commission); 24 C.F.R. pt. 214 (2024) (HUD housing counseling program).

[34] *See* 47 U.S.C. § 254 (2024) (universal service provisions added by the Telecommunications Act of 1996); 47 C.F.R. pt. 54, subpt. E (2024) (Lifeline program subsidizing telecommunications service for low-income consumers).

Wainwright's right to counsel—a protection against state power that the state itself provides.[35]

B. Permitting use, foreclosing abuse

Anonymity has always been double-edged. Bitcoin enabled private exchange, and it enabled the Silk Road marketplace and its successors.[36] End-to-end encryption protects the dissident and the predator both.[37] The history of privacy-enhancing technology is, in significant part, a history of regulators negotiating how much abuse to tolerate in exchange for how much legitimate protection. Agent-mediated commerce will replay that negotiation at a higher level of abstraction.

The distinctive difficulty is that agent-mediated anonymity is not a shield around one dimension of the transaction—value, content, identity—but around the transaction as a whole. A regulator responding to Silk Road could focus on the payment rail: Bitcoin in, fiat out, KYC at the on-ramps and off-ramps. A regulator responding to encrypted messaging could focus on the endpoint: the device, the app provider, the metadata. A regulator responding to agent-mediated commerce faces a technology whose function is to make every point in the transaction opaque simultaneously. There is no clean rail at which identification can be mandated without collapsing the very anonymity the regime is meant to preserve.

Regulation in an anonymity regime may need to operate somewhere other than the transaction layer itself.[38] Two directions are worth considering. *Above* the transaction, at the gateway: agent providers might bear obligations to maintain audit trails releasable only under judicial

---

[35] *See* Gideon v. Wainwright, 372 U.S. 335, 344 (1963) ("reason and reflection require us to recognize that in our adversary system of criminal justice, any person haled into court, who is too poor to hire a lawyer, cannot be assured a fair trial unless counsel is provided for him").

[36] United States v. Ulbricht, 858 F.3d 71 (2d Cir. 2017); *see also* Nicolas Christin, *Traveling the Silk Road: A Measurement Analysis of a Large Anonymous Online Marketplace*, in Proc. 22d Int'l Conf. on World Wide Web 213, 220–22 (2013) (estimating Silk Road's annual revenue and transaction volume).

[37] *See* Steven Levy, Crypto: How the Code Rebels Beat the Government—Saving Privacy in the Digital Age (2001).

[38] The literature on cryptocurrency regulation, KYC/AML obligations, and platform governance has generally assumed that the transaction remains the operative regulatory unit. *See, e.g.*, Sarah Jane Hughes & Stephen T. Middlebrook, *Advancing a Framework for Regulating Cryptocurrency Payments Intermediaries*, 32 Yale J. on Reg. 495 (2015); Kate Klonick, *The New Governors: The People, Rules, and Processes Governing Online Speech*, 131 Harv. L. Rev. 1598 (2018). By making every point in the transaction opaque simultaneously, agent-mediated commerce forecloses that assumption and requires the regulatory architecture to sit above or below the transaction layer rather than within it.

authorization, preserving a pathway to identification when circumstances warrant. *Below* the transaction, at the remedial layer: victims of agent-enabled fraud might be made whole through ex ante insurance or distributive compensation rather than retrospective investigation that anonymity renders impractical.

The KYC and anti-money-laundering regimes illustrate the challenge. Existing financial regulation requires identification of counterparties to certain transactions as a condition of participation in the banking system.[39] AI agents transacting in cryptocurrency or single-use card numbers sit uneasily with these requirements. A response that required the agent itself to be identified would defeat the architecture; one that exempted agent-mediated transactions entirely would turn the agent into a laundering tool.

C. The retailer ban

The private side of the market is likely to be less accommodating. Retailers have strong incentives to identify their customers. Identification enables personalized pricing, targeted upselling, profile building, resale of behavioral data, and—not least—the commercial surveillance and discrimination that Part I described as socially costly but that are, from the retailer's point of view, profit-generating. A retailer facing a customer who presents as an agent sees a customer whose transactional value has been narrowed to the margin on the immediate sale. Everything else the retailer would normally extract is gone.

The natural response is refusal. A retailer can amend its Terms of Service to ban agent-mediated transactions, require proof of human identity before completing a purchase, implement CAPTCHA-like challenges designed to defeat automation, or simply decline to serve any account

[39]Bank Secrecy Act, 31 U.S.C. § 5318(l) (2024) (requiring financial institutions to implement customer identification programs); 31 C.F.R. § 1020.220 (2024) (bank Customer Identification Program rule requiring identity-verification procedures sufficient to form a reasonable belief that the bank knows the customer's true identity); Fin. Action Task Force, *International Standards on Combating Money Laundering and the Financing of Terrorism & Proliferation: The FATF Recommendations*, Recommendation 10 (2012, updated Oct. 2025), https://www.fatf-gafi.org/en/publications/Fatfrecommendations/Fatf-recommendations.html; Fin. Crimes Enf't Network, *Application of FinCEN's Regulations to Certain Business Models Involving Convertible Virtual Currencies*, FIN-2019-G001, at 15–17 (May 9, 2019), https://www.fincen.gov/sites/default/files/2019-05/FinCEN/%20Guidance/%20CVC/%20FINAL/%20508.pdf./label%7Bfn:kyc}

it suspects of being an agent. Each restores the retailer's profile-building capacity by forcing the customer to present as a human with a persistent identity. Under current law, each is presumptively permissible: private actors generally set the terms of their own commerce, and the idea that a retailer might be *obligated* to serve anonymous customers cuts sharply against the background presumption of freedom of contract.[40]

But freedom of contract is not absolute. Public-accommodations law already limits a retailer's freedom to refuse service on protected-class grounds.[41] If agent-mediated anonymity is, as Part I argued, an effective protection against algorithmic discrimination, then a retailer's refusal to serve agents is a refusal to serve customers through a channel that prevents discrimination. The functional consequence is the restoration of the retailer's capacity to discriminate. The law's background presumption—that private actors set their own terms—runs up against the law's foreground commitment—that private actors may not set those terms in ways that enable discrimination.

Framed narrowly, that is a business decision—a private actor setting the terms of its own commerce. Framed more honestly, it is a demand for identification as a condition of participation in the commercial sphere. Forced identification has its own history in American civil rights law, and that history is not on the side of the demand. Pass laws in colonial and antebellum America required free Black people to carry papers establishing their right to be where they were.[42] Voter identification debates continue to turn on the disparate impact of identification requirements on the same populations the civil rights regime is designed to protect.[43] A commercial requirement

---

[40] *See generally* P.S. Atiyah, The Rise and Fall of Freedom of Contract (1979) (tracing the rise and decline of contractual autonomy as an organizing principle of Anglo-American private law); Lawrence M. Friedman, Contract Law in America: A Social and Economic Case Study (1965) (examining the social and economic forces shaping American contract doctrine).

[41] Civil Rights Act of 1964, tit. II, 42 U.S.C. § 2000a (2024); *see* Heart of Atlanta Motel, Inc. v. United States, 379 U.S. 241 (1964); *cf.* Masterpiece Cakeshop, Ltd. v. Colo. Civil Rights Comm'n, 584 U.S. 617 (2018) (addressing the tension between public-accommodations law and religious and free-speech objections to serving certain customers).

[42] *See* Ira Berlin, Slaves Without Masters: The Free Negro in the Antebellum South 92–96, 327–40 (1974) (describing registration requirements, certificates of freedom, and pass laws imposed on free Black people throughout the antebellum South); Thomas D. Morris, Southern Slavery and the Law, 1619–1860, at 22–36 (1996) (tracing the legal architecture of racial control, including documentation requirements distinguishing free Black people from enslaved people).

[43] *See* Veasey v. Abbott, 830 F.3d 216, 249–50 (5th Cir. 2016) (en banc) (holding that Texas's voter identification law had

that every purchaser be identifiable is structurally analogous to these earlier regimes.[44]

Conclusion

Orwell's Winston Smith thought the inside of his skull was the last uncolonized ground. This Article has argued that the perimeter of the private can, in at least one dimension of modern life, be pushed back outward again. AI agents can return to commerce a degree of anonymity practically unavailable since cash. The civil rights implications are unusual: the mechanism of algorithmic discrimination collapses when the counterparty is a shell, as does civil rights law's detection-and-remedy architecture. Anonymity protects on one side and forecloses on the other, falling hardest on the populations the civil rights project was built to protect. Whether the law recognizes the possibility, and whether it protects the recognition against private actors incentivized to deny it, is the question next-generation commercial regulation must face.

a discriminatory effect on minority voters in violation of Section 2 of the Voting Rights Act); Crawford v. Marion Cnty. Election Bd., 553 U.S. 181, 198–203 (2008) (upholding Indiana's voter-ID law but acknowledging the burden on voters who lack qualifying identification); Zoltan L. Hajnal, Nazita Lajevardi & Lindsay Nielson, *Voter Identification Laws and the Suppression of Minority Votes*, 79 J. Pol. 363, 371–73 (2017).

[44]The structural analogy between a retailer's Terms-of-Service ban on agent-mediated transactions and historical forced-identification regimes (pass laws, voter ID requirements) appears novel to AI commerce. Prior scholarship on algorithmic discrimination has generally focused on how systems use observed or inferred traits once individuals are legible to the system, *see, e.g.*, Kim, *supra* note 18, at 860–61; Anupam Chander, *The Racist Algorithm?*, 115 Mich. L. Rev. 1023, 1028–31 (2017), rather than on whether forced identification is itself a contested civil rights question.